

**PATENT APPLICATION
ATTORNEY DOCKET NO. KON03-0004**

5

10

**METHOD AND APPARATUS FOR
OPTIMIZING CONTENT DELIVERY ON
LOCAL SUBNETS**

15

Inventors: Wade L. Hennessey and John B. Wainwright

BACKGROUND

20

Field of the Invention

[0001] The present invention relates to systems that communicate across computer networks. More specifically, the present invention relates to a method and apparatus for optimizing content delivery on local subnets.

25

Related Art

30

[0002] In the recent past, the majority of local area networks were shared carrier networks. These shared carrier networks were typically unswitched, half-duplex networks. If one node on the network wanted to download a piece of content while all of the other nodes were inactive, the one node enjoyed virtually 100% of the available network bandwidth. However, if the one node was

competing with another node on the local subnet for bandwidth, the bandwidth would theoretically be split between the two nodes, and the one node would see a transfer rate of about 50% of the previous transfer rate. The problem, however, is worse than this because each additional node on the network requires an
5 increasing amount of network overhead to handle collision detection and collision avoidance.

[0003] With the emergence of low cost switches and increasing Ethernet speeds, these problems seem to have disappeared. On a switched network, traffic is routed to only the portions of the network where it is intended. Thus, two nodes
10 on a switched network transferring files between themselves typically do not affect other clients on the same subnet.

[0004] Recent times, however, have brought new wireless networking technologies to the market place. While the majority of wireless networks provide a large number of benefits, they have at least one major drawback. Most wireless
15 networks are shared carrier networks. This means that the bandwidth that each nodes consumes on the network decreases the bandwidth available to other nodes by at least that amount. While operating in “infrastructure mode,” when a first node on a wireless network communicates with a second node, the first node typically transmits data to an access point which relays the data to the second
20 node. Thus, when two nodes are communicating on a wireless network in this way, they are able to use less than half of the theoretical capacity of the network. Add a handful of additional nodes to the mix, and it is easy to see how quickly the network performance deteriorates.

[0005] Businesses, which have been some of the earliest adopters of
25 wireless networks, often encounter a situation where many clients on the same subnet require the same content. For example, when a company releases a new presentation or a software update, that content is often distributed to virtually all

of the nodes on the network. Moreover, businesses often employ a subscription-based delivery service in which every node that will eventually need a specific item of content has a subscription to a category that the content belongs. Each node checks at random intervals to see if there are new pieces of content for their subscription categories. After a new item of content becomes available, each node will eventually perform its content check and will subsequently download the new item of content. In this way, each node will independently download the same item of content, which consumes a large amount of wireless networking bandwidth. For example, if there are 10 nodes on a wireless subnet, trying to access the same 100 MB file, the 100MB may have to be downloaded 10 separate times across the wireless network. From this example it is easy to see how quickly the wireless subnet can be overwhelmed with traffic.

[0006] Hence, what is needed is a method and an apparatus for distributing content to clients on a local subnet in a manner that minimizes the problems described above.

SUMMARY

[0007] One embodiment of the present invention provides a system that facilitates optimizing content delivery on a network. During operation, the system receives an item of content at a first client. During the download of the content or after downloading the content, the first client receives a broadcast request for the content from a second client on the same local subnet. Upon receiving the broadcast request, the first client sends a broadcast response to the local subnet, wherein the broadcast response identifies a multicast address to which the first client will deliver the content. The first client then delivers the content to the multicast address so that the second client and any other interested clients on the local subnet can receive the content.

[0008] In a variation on this embodiment, receiving the content at the first client involves first sending a broadcast to the local subnet requesting the content.

If a response to the broadcast is received, the first client receives the content via a multicast transmission from another client on the local subnet. On the other hand,
5 if a response to the broadcast is not received, the first client receives the content from a source outside of the local subnet via a unicast transmission.

[0009] In a further variation, if the first client receives the content via a unicast transmission, the first client additionally sends a second broadcast to the local subnet announcing that the first client is receiving the content, thereby
10 allowing other clients on the local subnet to request to receive the content from the first client.

[0010] In a further variation, the second broadcast contains information about the content, including subscription information. This subscription information allows other clients to determine if they have a subscription to receive
15 the content and should therefore request to receive the content.

[0011] In a further variation, the first client sends the second broadcast to the local subnet prior to receiving all of the content, whereby the first client can start transferring the content to other clients on the local subnet via multicast prior to receiving all of the content.

20 **[0012]** In a variation on this embodiment, the first client receives a broadcast message from another client on the local subnet announcing that the other client is transmitting another item of content, and including a second multicast address for delivery of the other content. The first client then determines if it needs the other content, and if so, receives the other content via
25 the second multicast address.

[0013] In a variation on this embodiment, the local subnet is a shared-carrier network.

[0014] In another variation, the local subnet is a wireless network.

[0015] In another variation, the wireless network adheres to the 802.11x protocols.

[0016] In a variation on this embodiment, the second client starts
5 receiving the multicast of the content while the multicast is already in progress.

BRIEF DESCRIPTION OF THE FIGURES

[0017] FIG. 1 illustrates a local area computer network in accordance with an embodiment of the present invention.

10 [0018] FIG. 2 presents a flowchart illustrating the process of optimizing content distribution on a local subnet in accordance with an embodiment of the present invention.

[0019] FIG. 3 presents a flowchart illustrating the process of acquiring content in accordance with an embodiment of the present invention.

15

DETAILED DESCRIPTION

[0020] The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed
20 embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features
25 disclosed herein.

[0021] The data structures and code described in this detailed description are typically stored on a computer readable storage medium, which may be any

device or medium that can store code and/or data for use by a computer system. This includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital versatile discs or digital video discs), and computer instruction signals embodied in a transmission medium (with or without a carrier wave upon which the signals are modulated). For example, the transmission medium may include a communications network, such as the Internet.

Local Area Network

10 [0022] FIG. 1 illustrates a local area computer network 100 in accordance with an embodiment of the present invention. Network 100 can generally include any type of wire or wireless communication channel capable of coupling together computing nodes. Coupled to network 100 are clients 102-116 and server 120. Clients 102-116 can generally include any node on network 100 including
15 computational capability and including a mechanism for communicating across network 100. Server 120 can generally include any computational node including a mechanism for servicing requests from a client for computational and/or data storage resources.

20 [0023] In one embodiment of the present invention, network 100 is a shared carrier wired network that operates in unswitched, half-duplex mode. In another embodiment of the present invention, network 100 is a wireless network that adheres to one of the 802.11x protocols. In either of these embodiments, when client 102 communicates with client 104, packets communicated between clients 102 and 104 tie up the shared carrier or wireless frequency, reducing the
25 bandwidth available to the other clients on network 100. For the purposes of the present invention, clients 102-116 are all located on local subnet 101, while server 120 is not on local subnet 101.

Process of Optimizing Content Distribution on a Local Subnet

[0024] FIG. 2 presents a flowchart illustrating the process of optimizing content distribution on local subnet 101 in accordance with an embodiment of the present invention. The system starts when an item of content is received at client 102 (step 202). Next, client 102 receives a broadcast request for the content from client 104 on network 100 (step 204). Note that client 102 may receive the broadcast from client 104 while client 102 is still receiving the content, as well as after client 102 has received the content. In response to this request, if client 102 still has the content, client 102 sends a broadcast to local subnet 101 on network 100 identifying the content and any associated subscription and specifying a multicast address to which client 102 will send the content (step 206). Client 102 then delivers the content to the multicast address (step 208).

[0025] By sending the broadcast to local subnet 101 that identifies the content and the multicast address, client 102 allows other clients on local subnet 101 on network 100 to determine if they need the content, and if so, to receive the content without consuming any additional bandwidth on network 100. In a subscription-based environment, every client that has a subscription to the source of the content, but does not yet have this particular item of content, can respond to the broadcast and can retrieve the content via client 102's multicast of the content.

[0026] By utilizing this system in a subscription-based environment (where clients check for new items in their subscriptions at a random interval over a certain period of time) it is possible to deliver the content to an arbitrary number of clients on local subnet 101 with the actual bits being transmitted only twice. The first transfer loads the content from the source to the original client, for example client 102, and the second transfer multicasts the content from client 102 to all of the other interested clients on local subnet 101. Note that client 102

could alternatively broadcast the content to the whole subnet, but this broadcast would cause other clients on local subnet 101 to have to examine each broadcast packet even if the other clients are not interested in them, thus causing additional work for the processor on each client. Note that although broadcast is not as efficient as multicast, it is still an improvement over multiple unicast downloads by each client. By using multicast, clients that do not want the content do not configure themselves to receive the content on the multicast address, and are thereby able to filter out the packets at the network interface level, without having to interrupt the processor.

10

Initially Acquiring the Content

[0027] FIG. 3 presents a flowchart illustrating the process of acquiring an item of content in accordance with an embodiment of the present invention.

When client 102 on network 100 initially wants an item of content, client 102 first sends out a broadcast on local subnet 101 to see if any other client on local subnet 101 already has the content (step 302). Next, client 102 determines if there is a response to the request (step 304). If so, client 102 receives the content via a multicast transmission from the other client (step 306). If not, client 102 receives the content via a unicast transmission from the content's original source, such as server 120, from outside local subnet 101 (step 308).

[0028] The foregoing descriptions of embodiments of the present invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.